



(19) **United States**

(12) **Patent Application Publication**
Rao

(10) **Pub. No.: US 2005/0086170 A1**

(43) **Pub. Date: Apr. 21, 2005**

(54) **SYSTEM AND METHOD FOR PROCESSING
PARTIALLY UNSTRUCTURED DATA**

Publication Classification

(76) **Inventor: Srinivasan N. Rao, New York, NY
(US)**

(51) **Int. Cl.7 H04K 1/00**

(52) **U.S. Cl. 705/50**

Correspondence Address:
**DOCKET ADMINISTRATOR
LOWENSTEIN SANDLER PC
65 LIVINGSTON AVENUE
ROSELAND, NJ 07068 (US)**

(57) **ABSTRACT**

A system and method for processing partially unstructured data relating to a financial security. The system and method resolve first- and second-identifying data from the partially unstructured data and determine whether a security is defined by the first-identifying data and the second-identifying data. Additionally, the system and method resolve trade information relating to the security identifier from the partially unstructured data. If a security is defined by the resolved identifying data, a security identifier representing the defined security, along with the trade information relating to the defined security, are output.

(21) **Appl. No.: 10/740,058**

(22) **Filed: Dec. 18, 2003**

Related U.S. Application Data

(60) **Provisional application No. 60/511,591, filed on Oct.
15, 2003.**

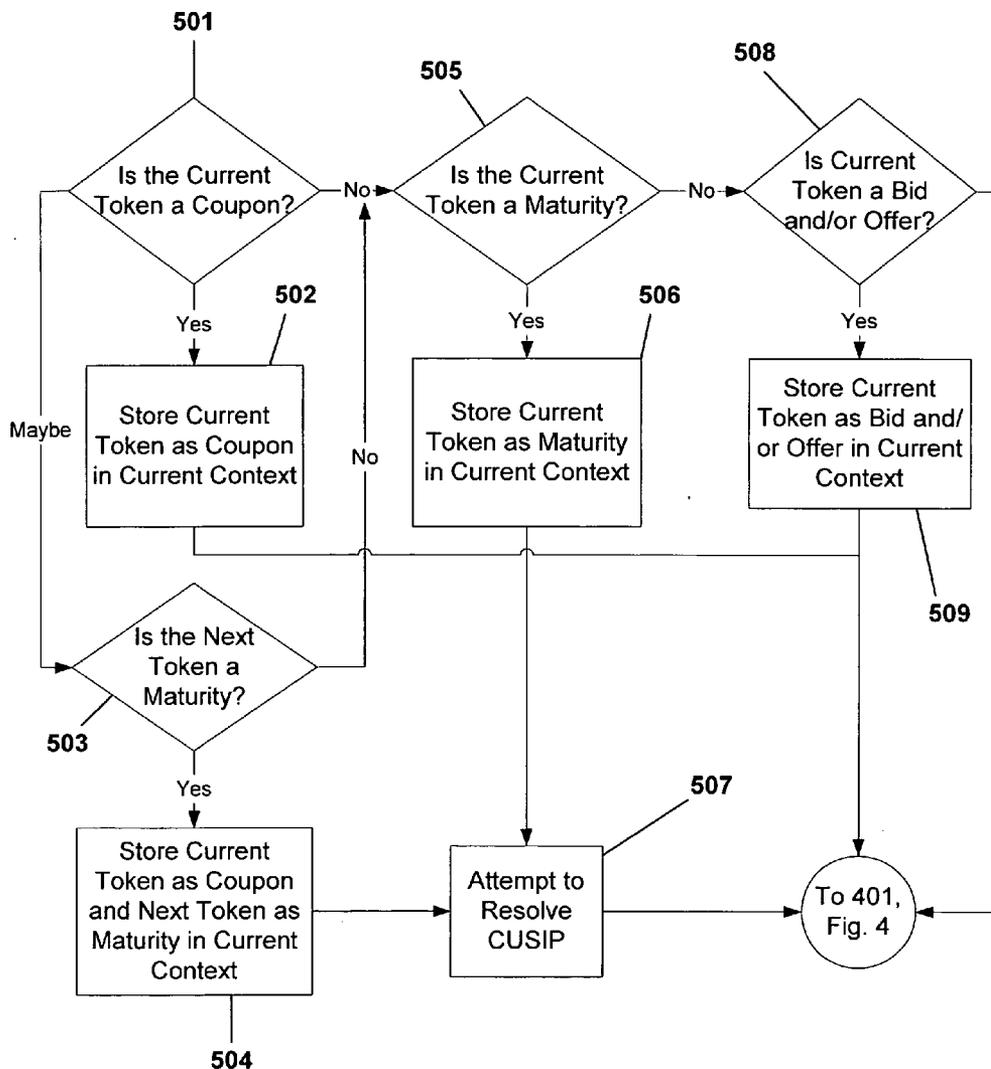


Fig. 1

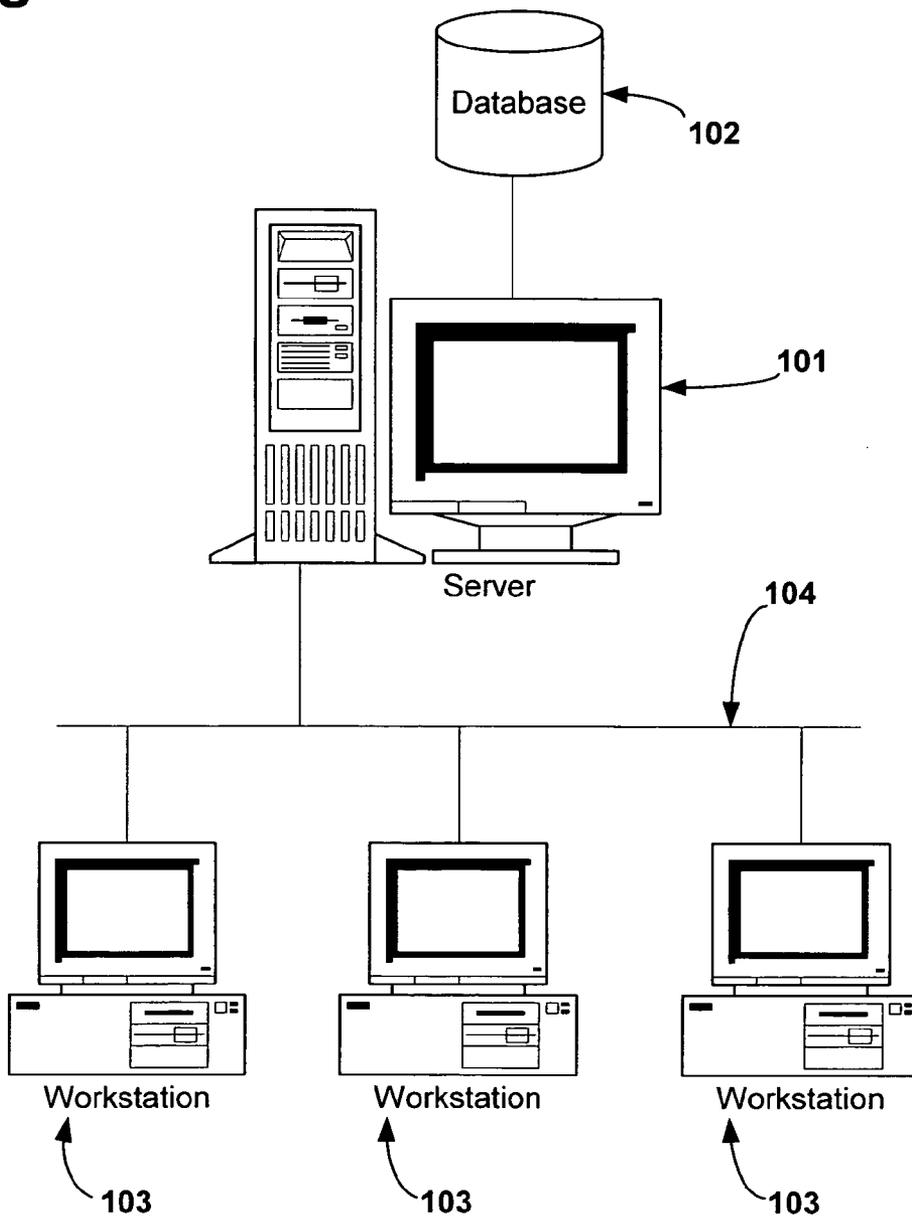


Fig. 2

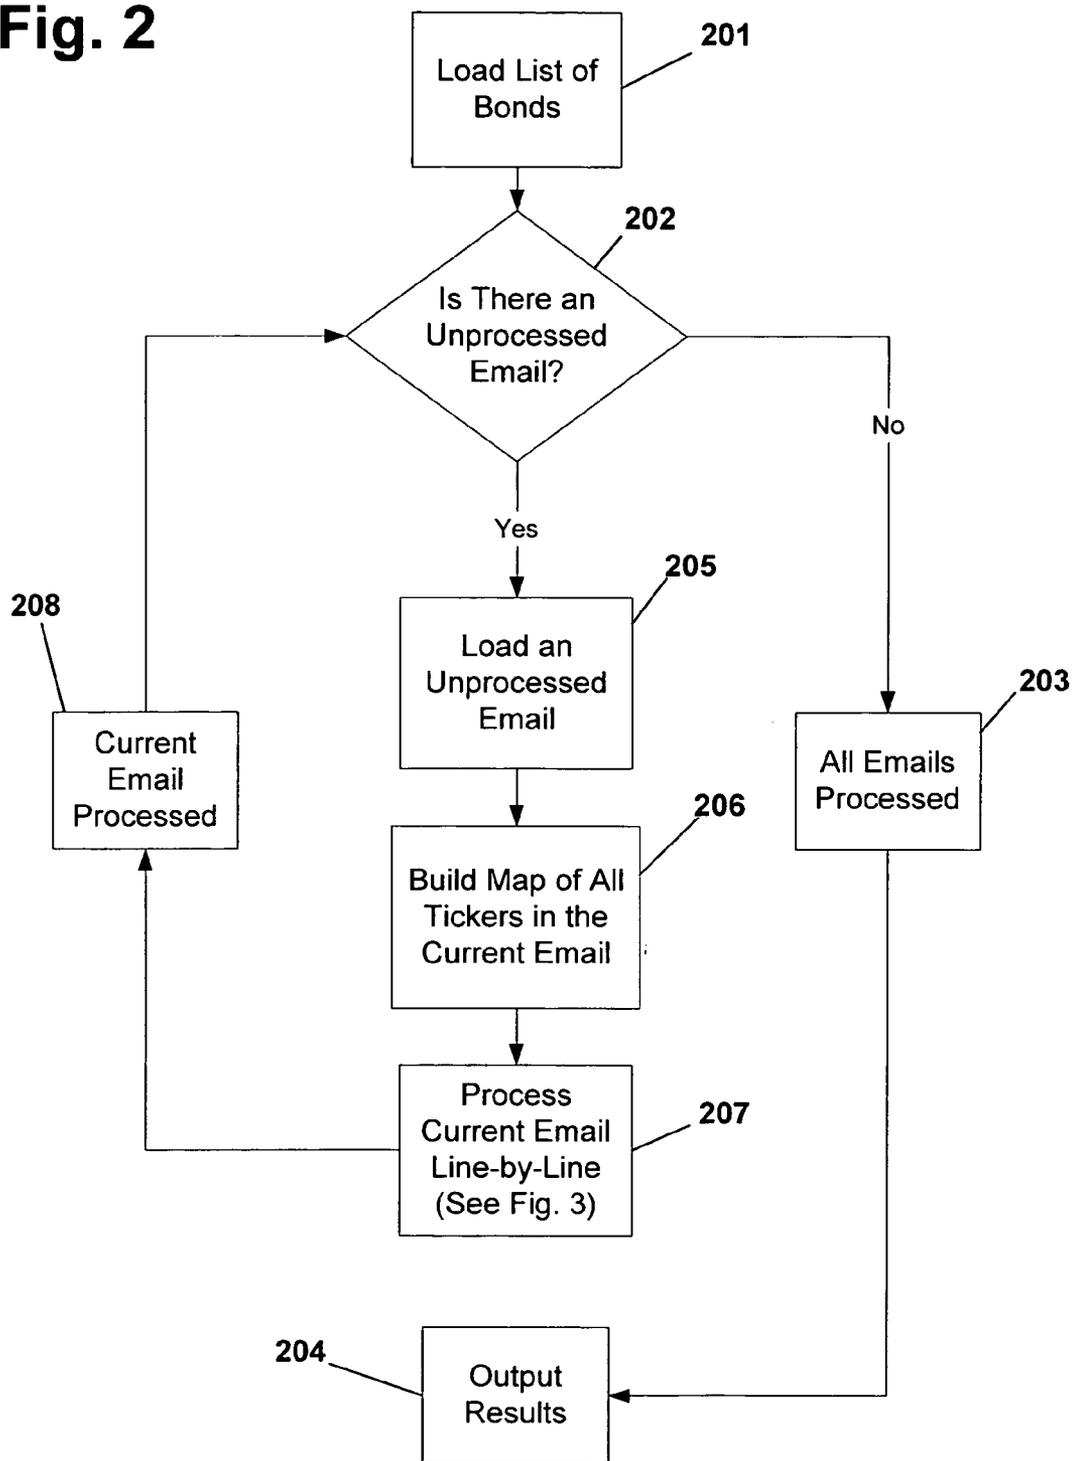


Fig. 3

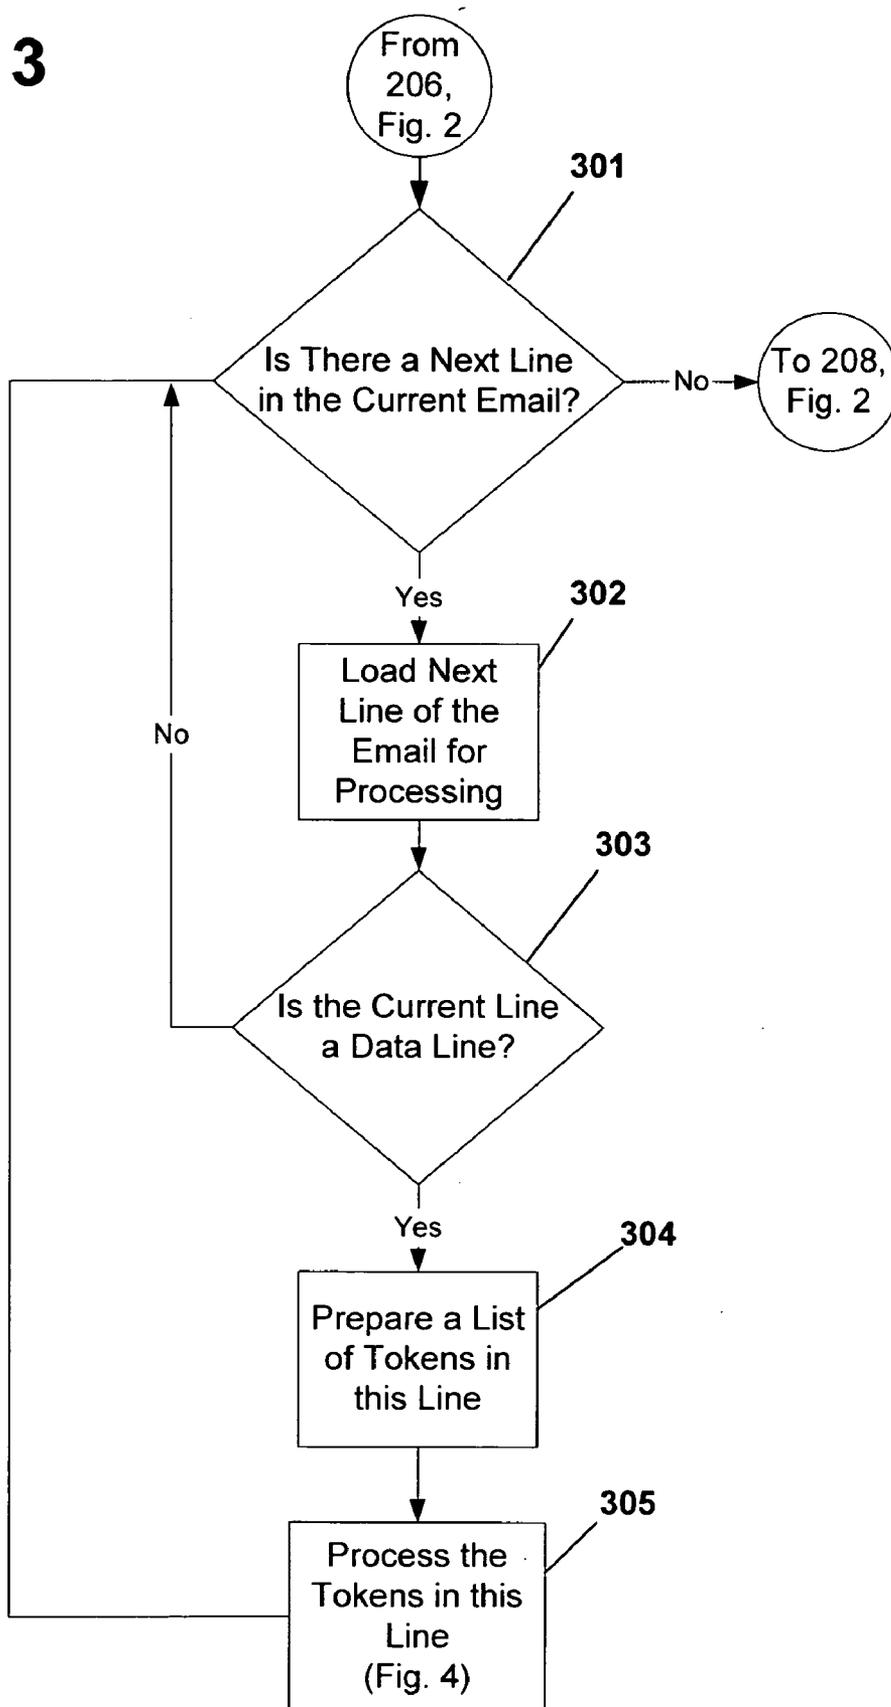


Fig. 4

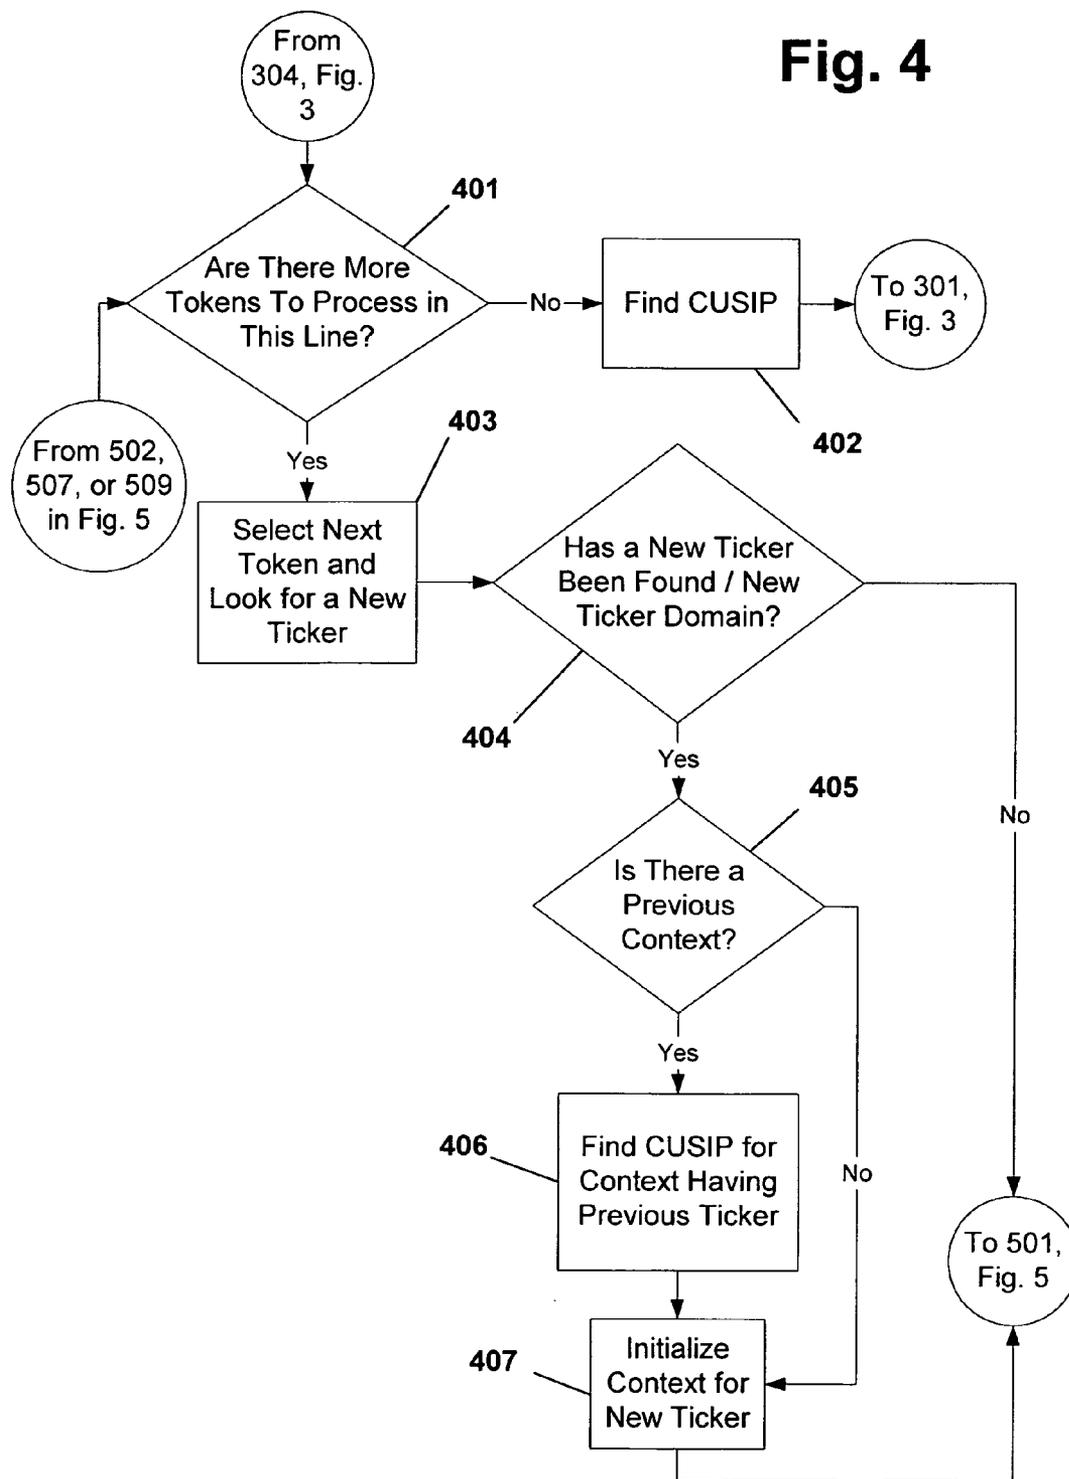
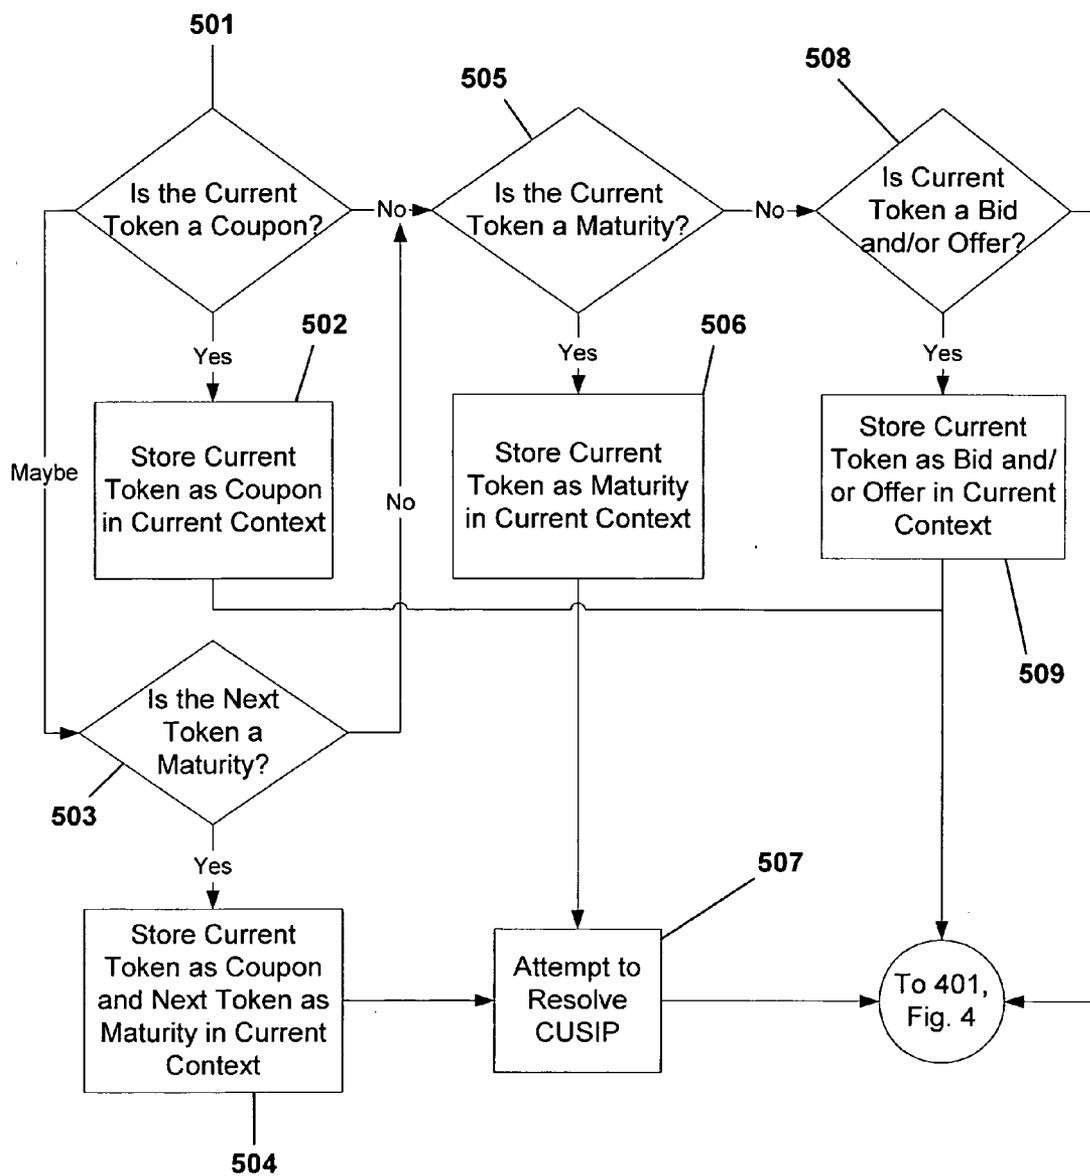


Fig. 5



SYSTEM AND METHOD FOR PROCESSING PARTIALLY UNSTRUCTURED DATA

CROSS-REFERENCE TO RELATED APPLICATION

[0001] This application claims the benefit of U.S. Provisional Application No. 60/511,591, filed Oct. 15, 2003, which is hereby incorporated herein by reference.

REFERENCE TO COMPUTER PROGRAM LISTING APPENDIX

[0002] The file of this patent application includes a computer program listing appendix stored on two identical read-only Compact Discs. Each Compact Disc has the computer program listing appendix stored as a file named "appendix1.doc" that was created on Nov. 7, 2003 and is 160,768 bytes in size. This computer program listing appendix is hereby incorporated herein by reference.

FIELD OF THE INVENTION

[0003] This invention relates to a system and method for processing partially unstructured data to extract valuable information from the partially unstructured data. In particular, this invention relates to processing partially unstructured data, such as text, to extract information of interest, such as information relating to the trading of securities. This invention enables traders of securities to access a higher quantity of trade information than they would ordinarily be able to access.

BACKGROUND OF THE INVENTION

[0004] For the trader of securities, it is very important to know what the best available prices are on the street in a timely manner and to be able to use the trading opportunities that these prices present before the window of opportunity closes. Nowhere is this more important than for bond trading. Typically, a Credit Default Swap trader receives information about bond prices in the form of emails. The bulk of these emails arrive within a very short period of time around the time when the markets open, and the information contained within these emails is valuable only for a limited period of time. It is common for traders to receive hundreds of these emails in the morning. Buried within these emails are often good trading opportunities.

[0005] In the conventional arrangement, the trader had to manually read through each of these emails to find out what the prevailing bond prices are being offered on the street. However, the trader often cannot read through all of these emails before the window of opportunity closes for taking advantage of the information in these emails. For every email the trader does not have time to read, he or she misses an opportunity to earn a profit.

[0006] Further, no rigid formatting convention for these types of emails exists. They are fairly unstructured and often differ significantly from one-another. For example, an email may have lines talking about an impending vacation and then may have lines stating, "by the way, I want to sell this particular bond at this particular price." Also, the email may or may not provide all of the information commonly used to identify a particular bond. Therefore, lack of consistent formatting in emails presents a technical problem for

extracting trading opportunity information from such emails with a relatively high rate of success.

SUMMARY OF THE INVENTION

[0007] These problems are addressed and a technical solution achieved in the art by this invention, which provides a system and method for processing partially unstructured data relating to financial securities. In particular, this system and method resolve first-identifying data from the partially unstructured data, resolve second-identifying data from the partially unstructured data, and determine whether a security is defined by the first-identifying data and the second-identifying data when the second-identifying data is of a predetermined type. The system and method also resolve third-identifying data from the partially unstructured data and determine whether a security is defined by the first-identifying data, the second-identifying data, and the third-identifying data. Additionally, the system and method resolve trade information relating to the security identifier from the partially unstructured data. If a security is defined by the first- and second-identifying data, or by the first-, second-, and third-identifying data, a security identifier representing the defined security is output along with the trade information relating to the security. Optionally, it is determined whether a security is unambiguously defined by the identifying data. In one embodiment, the first-identifying data represents a ticker, the second-identifying data represents a coupon or a maturity, the third-identifying data represents the other of a coupon or a maturity that the second-identifying data represents, and the predetermined type is a maturity.

[0008] Described in a different manner, the system and method identify at least one of a plurality of predefined data vectors from partially unstructured data. The partially unstructured data includes a plurality of data items having positions relative to each other in the partially unstructured data. The system and method determine a position of each of one or more data items of a first type from the plurality of data items in the partially unstructured data. A data item of a second type is selected from the plurality of data items in the partially unstructured data. The system and method also select one of the one or more data items of the first type based on its position relative to the selected data item of the second type. A data item of a third type and a data item of a fourth type are selected from the plurality of data items in the partially unstructured data. The system and method identify a predefined data vector from the plurality of predefined data vectors from the selected data item of the first type, the selected data item of the second type, and the selected data item of the third type. The data item of the fourth type and an identifier representing the identified data vector are output. Examples of data items of a first, second, third, and fourth type are a ticker, coupon, maturity, and trade information, respectively. Alternate examples of data items of a first, second, third, and fourth type are a ticker, maturity, coupon, and trade information, respectively. An example of an identifier is a CUSIP. The data items of the first, second, third, and fourth type, along with the identifier, may be stored in a context.

[0009] This invention provides a technical solution in that it processes the vast quantity of emails that a trader receives in the morning, and extracts from many of them, the identities of the securities, such as stocks and/or bonds, and

trade information relating to each of the identified securities, such as bid and/or offer prices. The extracted information is then accessible to the trader in the morning when the markets open, the time period when it is needed. The invention provides much more information regarding prevailing bond prices than would normally be available if the trader has to manually read through each of the emails.

BRIEF DESCRIPTION OF THE DRAWINGS

[0010] A more complete understanding of this invention may be obtained from a consideration of this specification taken in conjunction with the drawings, in which:

[0011] FIG. 1 is an example of a hardware arrangement implementing the preferred embodiment; and

[0012] FIGS. 2-5 are flowcharts depicting the major processing steps performed by the preferred embodiment.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT OF THE INVENTION

[0013] I. Definitions:

[0014] Prior to discussing the details of the preferred embodiment, several definitions of terms used throughout this specification are set forth below.

[0015] 1) Ticker: a system of letters used to uniquely identify a stock or mutual fund.

[0016] 2) Coupon: the interest rate stated on a bond when it's issued. Also referred to as "Rate."

[0017] 3) Maturity: The length of time until the principal amount of a bond must be repaid.

[0018] 4) Bid: Price at which to buy a security. A bid is considered a type of trade information relating to the security.

[0019] 5) Offer: Price at which to sell a security. Also referred to as "Ask." An offer is considered a type of trade information relating to the security.

[0020] 6) Token: a segment of data that can represent one of (1) a security's coupon, (2) maturity, or (3) bid and/or offer.

[0021] 7) CUSIP Number or CUSIP: A number used to identify all U.S. and Canadian stocks and registered bonds. ("CUSIP" is a registered trademark of the American Bankers Association.) A security's CUSIP can be identified by its ticker, coupon, and/or maturity. Therefore, a ticker, coupon, and maturity are types of identifying information used to identify a CUSIP for a particular security. One having ordinary skill in the art will appreciate that a CUSIP could be represented as a data vector comprising the particular ticker, coupon, and maturity associated with the CUSIP in question as data items.

[0022] 8) CINS Number or CINS: A number used to identify all international stocks and registered bonds. A security's CINS can be identified by its ticker, coupon, and/or maturity.

[0023] 9) Ticker Domain: a region in an email that is associated with a particular ticker identified in the email, wherein if a token is located in this region, it is associated with the particular ticker.

[0024] 10) Context: A set of information relating to a particular security, the information including identifying data, such as the security's ticker, coupon, maturity, and bid and/or offer prices, wherein the identifying data can be used, among other things, to identify one or more CUSIP numbers that correspond to the identifying data.

[0025] II. Description:

[0026] The preferred embodiment of this invention is described in the context of processing emails containing information relating to bonds, wherein the bonds are identified by their CUSIP number. However, one having ordinary skill in the relevant art will appreciate that the disclosed system and method can be readily adapted to process data transmitted in different manners besides email. For instance, the data can be in the form of a regular text file, an image file that has been converted to a text file, or any type of file that can be parsed by a computer to extract text information. One having ordinary skill in the relevant art will also appreciate that the disclosed system and method can be readily adapted to process data besides bonds, including other security types, such as stocks and/or mutual funds. Additionally, the disclosed system and method can be readily adapted to search for other types of identifiers besides CUSIP numbers, such as CINS numbers, or any other means to identify data, without departing from the scope of this invention.

[0027] Prior to discussing the details of the preferred embodiment, an example of a portion of an email received by a trader will be explained. Consider the following example, shown in Table I below, of an excerpt from an email received by a trader.

TABLE I

IBM	07/05	8	100/
	04/07	10	/230

[0028] In Table I, the first line refers to a single security. The letters "IBM" refer to the ticker relating to the security. "07/05" refers to a maturity month and year of the bond, "8" refers to the coupon, or rate of the security, and the "100/" is the bid price because it is followed by a "/". The second line refers to another security with the same ticker. The "04/07" refers to the security's maturity month and year, the "10" refers to the coupon, and the "/230" refers to the offer price because it is preceded by a "/".

[0029] Needless to say, most emails are not this structured, but contain the same or similar types of information. The details of how the preferred embodiment of the invention processes all of these emails, whether structured or not, will not be set forth, beginning with reference to FIG. 1.

[0030] FIG. 1 depicts a preferred hardware arrangement implementing the present invention. In FIG. 1, a server computer 101, either containing a database 102, or being in communication with a database 102, is in communication via communication mechanism 104 with one or more workstation computers 103. Any method of communicating between computers may be used between the server 101 and the workstations 103, and the server 101 and the database 102, if not contained within the server 101. The communication mechanism 104 need not be a hardwired network, and

may be wireless, or a combination of both. Workstations **103** do not have to be actual desktop computers, as shown in **FIG. 1**, and can be other types of computers, such as laptops, hand-held devices, or any device that includes a computer.

[**0031**] In the preferred embodiment, the database **102** stores all of the emails received from clients, typically via the Internet, and it also stores a list of all bonds, including their ticker, coupon, maturity, and CUSIP number. Traders have access to workstations **103**, where they can login to access their particular account. Logging in includes communication with the server **101** to transmit that particular trader's information to the trader's workstation **103**. Also according to the preferred embodiment, the present invention is implemented as a program stored on the server **101**, where it is executed to process the received emails and extract the bond CUSIP numbers and their bid and/or offer prices. However, the program can be stored on one or all of the workstations **103**, and executed from any location. Also, it is possible to have the program and database all stored on a single computer.

[**0032**] The manner of processing the emails according to the preferred embodiment of this invention will now be described with reference to **FIGS. 2-5**. **FIG. 2** provides a high level view of the entire process performed by this embodiment. The subsequent figures, **FIGS. 3-5**, provide more detail regarding **207** shown in **FIG. 2**. With reference to **201** in **FIG. 2**, the list of bonds, containing each bond's ticker, coupon, maturity, and CUSIP number, is initially downloaded from the database **102** into the local memory of the computer performing the email processing, such as the server **101**. Next, it is determined whether or not any unprocessed emails exist in the database **102**. If none exist, it is determined that all of the emails have been processed at **203** and any bond CUSIPs and their corresponding bid and/or ask prices that have been identified through the email processing are stored in the database **102** and output via email to the traders at **204**.

[**0033**] If unprocessed emails remain in the database **102**, the next of those emails is downloaded from the database **102** and stored in local memory for processing at **205**. Some initial preprocessing of the downloaded email is performed at this time to eliminate the header of the email and, optionally, to store general statistical information about the email, such as storing the number of occurrences of the word "bid" and "offer" that are present in the email. The statistical information may be useful in identifying bid and/or ask prices included in the email.

[**0034**] Next, a map of all of the tickers in the current email is generated at **206**. The map stores each ticker name found in the email as well as its position in the email. This map will subsequently be used to determine which ticker a particular token belongs.

[**0035**] To identify a ticker, the preferred embodiment processes the email line-by-line. Before looking for tickers in a line, the line is preprocessed to correct formatting issues, such as making instances of "2x2" and "2x2" uniform. After the line has been preprocessed, the line is parsed one word at a time, comparing each word to the list of tickers provided by the downloaded bond list data **201**. If the word matches a ticker, several checks are executed to determine if it is in fact, not a ticker, even though it matches one in the list. In particular, if the word is preceded or succeeded by a "/" or

a "," followed by another word, such as "FON/AWE" or "FON,AWE", then the word is determined not to be a ticker, and it is skipped for the next word. Also, if it is a word like "cash bonds", "AT+344", "AT/+344", or "AT \$544" it is determined that the word is not a ticker, and it is ignored.

[**0036**] If the word that matches a ticker in the list is not eliminated by the above-described checks, the word is determined to be a ticker and its position in the line of the email is recorded. If the ticker is located at the start of a set of data, then the position of the ticker is not adjusted. An example of a ticker located at the start of a set of data is shown in Table II below, wherein "FON" is the ticker.

TABLE II

FON	6.25	11	360-370
-----	------	----	---------

[**0037**] If the ticker is located to the right of a start of a set of data, as shown for example in Table III below wherein "BA" is the ticker, then the position of the ticker is chosen to be the first word in the line that matches a word in the issuer's name for that ticker, i.e., "BOEING". Issuer names can be provided with the downloaded bond data **201**.

TABLE III

BOEING	CAPITAL	C(BA)
5.65	05/06	65-60

[**0038**] After the map of tickers is built, it is preferable to adjust the map such that if the left-most ticker in a line is not at position zero, then the positions of the tickers in that line are shifted to the left so that the left-most ticker in the line is at position zero. This simplifies subsequent processing.

[**0039**] At this point, a map of all tickers in the current email is generated, completing the processing described at **206** in **FIG. 2**. After this, the email is then processed line-by-line at **207** in an attempt to extract bond CUSIP numbers and the corresponding bid/offer prices. After all of the lines of the current email have been processed, it is determined that the current email has been completely processed at **208**, and the process repeats by checking the database **102** for a next unprocessed email at **202**.

[**0040**] Now the manner in which the email is processed line-by-line in an attempt to extract bond CUSIP numbers and the corresponding bid and/or offer prices at **207** will be described with reference to **FIG. 3**. At **301**, it is determined whether a next, unprocessed line in the email exists. If not, execution proceeds to **208** in **FIG. 2**, where it is decided that the current email has been completely processed. If an unprocessed line does exist, it is marked for processing at **302**. In other words, the unprocessed line is identified by a pointer, a corresponding array position, or loaded into a local variable, etc. The marked email line becomes the "current" email line for processing.

[**0041**] Next, it is determined whether the current email line is a data line at **303**. Data line means that the current email line has data that could represent a coupon, maturity, or a bid and/or offer. To determine if the current email line is a data line, this embodiment of this invention checks for data having the format of coupons, maturities, bids or offers.

For instance, the current email line must contain numbers to be a data line. Otherwise, no coupon, maturity, bid or offer is assumed to be present. Also, numbers having a “/” between them could be bid and offer. If numbers having the format of a coupon, maturity, bid, or offer are found, the current line is determined to be a data line and processing of the line continues at **304**. Otherwise, it is determined not to be a data line, and the current line is skipped. Execution then proceeds to **301** to check for a next unprocessed email line.

[**0042**] At **304**, a list of tokens in the line is prepared. For each token found, its position in the line is recorded. Preparing the list of tokens is achieved by searching the line for numbers and numbers separated by a “/” or a “-”. Numbers separated by a “/” or a “-” are considered a single token. Table IV below shows examples of tokens, wherein each row in the table represents a single token.

TABLE IV

05/06
5.65
10
100/
90-100

[**0043**] After a list of tokens has been prepared for the current line at **304**, the tokens in the line are processed to determine if they are coupons, maturities, or bids and/or offers at **305**. If tokens are identified as maturities, or if all of the tokens in a line have been processed, an attempt is made to identify one or more CUSIPs for the ticker and related token(s) that have been identified. This process is discussed in more detail below, with reference to **FIGS. 4 and 5**. However, before discussing this process, it is helpful to first define the usage of the terms “context” and “ticker domain”, which will be used throughout the remainder of this description.

[**0044**] A “context” is a set of stored information relating to a particular bond. This set of information includes identifying data, including the bond’s ticker, coupon, and maturity, which are used to attempt to resolve a CUSIP for the particular bond. An example of two contexts is shown in Table V below.

TABLE V

Context 1:	
Ticker:	BA
Coupon:	5.65
Maturity:	05/06
Bid Price:	100
Offer Price:	95
Context 2:	
Ticker:	BNI
Coupon:	Null
Maturity:	12/05
Bid Price:	104
Offer Price:	Null

[**0045**] Although Table V shows five data fields for ticker, coupon, maturity, bid price, and offer price, the context may include more than these data fields. When a token relating to a particular context is identified as a coupon, maturity, or bid price and/or offer price, the token’s data is then stored in the

corresponding field of the context. For example, if a current token pertaining to the ticker BNI has the data 6.375, and such data has been identified as a coupon, the null value for the coupon field in context 2 will be replaced with 6.375.

[**0046**] Each context relates to one of the tickers located in the email, as mapped at **206** in **FIG. 2**. It is possible however, to have more than one context relating to a single ticker in the situation where several sets of coupons and maturities are described with reference to a single ticker. When the context for a particular ticker is initialized, the data fields for coupon, maturity, bid price, and ask price are set to NULL. As data for these fields are extracted from the tokens in the email, their NULL values are replaced with the newly extracted data.

[**0047**] A “ticker domain” is a mechanism by which a token is associated with a particular ticker, and consequently, a particular context. This allows the data from the token to be placed in the appropriate context. For example, if an email contains the lines shown in Table VI below,

TABLE VI

BNI		
6.375	12/05	100-95
CSX		
7.25	05/04	105-100

[**0048**] the tokens 6.375, 12/05, and 100-95 are all in the BNI ticker domain, and their data will be stored in the context for the BNI ticker. The tokens 7.25, 05/04, and 105-100 are all in the CSX ticker domain, and will be stored in the corresponding context. The manner in which a token is associated with a ticker domain will be described later.

[**0049**] With a context and ticker domain defined, the processing of the tokens in a line will now be described with reference to **FIG. 4**, which is an exploded view of **305** in **FIG. 3**. The first action performed in processing the tokens in the current email line is to determine whether any unprocessed tokens exist in the current line at **401**. If all of the tokens in this line have been processed, an attempt is made to resolve a CUSIP for the current context **402**. The current context is the context for the ticker to which the previous token applied. In other words, if the previous token was a coupon for ticker “BA”, the current context is the context for ticker “BA”. It is noted that because the current line has been determined to be a data line, **303** in **FIG. 3**, at least one token exists in this data line, thereby preventing the scenario where the line has no token.

[**0050**] After attempting to resolve the CUSIP for the current context, the manner of which is explained in more detail later when discussing **FIG. 5**, the current line’s processing is complete, and the process returns to **301** in **FIG. 3**, where the email is checked for another unprocessed line. If more unprocessed tokens exist in the current line, the next unprocessed token is selected as the current token and a check is made to determine if the current token is in a new ticker domain at **403**. This is performed by searching for a ticker between the position of the previous token and the current token. If a ticker is found between the previous token’s position and the current token’s position, it is determined that the current token is in a new ticker domain.

If no new ticker is found, it is determined that the current token is in the previous ticker domain and the token data, if resolved, is added to the context pertaining to that ticker. If no new ticker is found, the process proceeds to **501** in **FIG. 5**.

[0051] If a new ticker has been found at **404**, it is determined that the current token refers to the new ticker, i.e., that it is in the new ticker's domain and a context for the new ticker should be initialized. Therefore, the previous context referring to the previous ticker will be processed. This begins with an initial check for whether a previous context exists at **405**, i.e., whether this newly found ticker is the first ticker in the email. If a previous context does not exist, i.e., this is the first ticker, then a first context is initialized for the new ticker at **407**. For example, if "BA" is the first ticker in the email, a first context will be initialized as shown in Table VII below.

TABLE VII

Context 1:	
Ticker:	BA
Coupon:	Null
Maturity:	Null
Bid Price:	Null
Ask Price:	Null

[0052] If a previous context does exist, i.e., there have been previous tickers, then an attempt is made to resolve a CUSIP for the context corresponding to the previous ticker at **406**, the process of which will be described later. After the attempt to resolve the CUSIP for the previous ticker has been made, a context is initialized for the new ticker at **407**.

[0053] Whether or not a previous context existed at **405**, the process ultimately moves to **501** in **FIG. 5**, wherein an attempt is made to resolve the current token.

[0054] The first step is to determine whether the current token is a coupon at **501**. This step is performed by analyzing the current token with respect to the current context. (Note that, although the following analysis is described in an order, such order is not necessarily required.) First, a check is made to find out if a coupon already exists in the current context, i.e., the coupon data field in the current context is not equal to null. If so, it is determined that the current token is not a coupon, and the processing moves on to **505** in **FIG. 5**. If a coupon does not exist in the current context, then the token is formatted to be in decimal form, if it is a fraction. This formatting simplifies subsequent data processing. Other formatting of the token may be performed to ensure that the token has the proper format of a coupon. Then, the formatted token is compared to the coupons in the list of bond data **201** to make certain that the formatted token has a value less than or equal to that of the maximum coupon value in the list of bond data. If the formatted token is greater than the maximum coupon value, it is determined that the current token is not a coupon, and processing proceeds to **505**.

[0055] If (1) the formatted token is less than or equal to the maximum coupon value in the list of bond data, (2) a maturity exists in the current context, and (3) if the current token cannot be a bid or an offer (discussed below), then it is determined that the current token is in fact a coupon. As such, the token is deemed to be resolved, it is stored in the

coupon field of the current context at **502**, and processing proceeds to the next token at **401** in **FIG. 4**. Otherwise, several more analyses are performed on the token before concluding that it is or is not a coupon.

[0056] If the current token, in its preformatted form, i.e., its original form, is a number with a fraction, such as "11 1/4", then the current token is determined to be a coupon, and is stored as such in the current context at **502** and processing proceeds to the next token at **401** in **FIG. 4**. If the current token is preceded by a single quote, a "/", a "-", or a "0", it is determined not to be a coupon, and processing proceeds to **505**.

[0057] If it is still undetermined whether or not the current token is a coupon, the preferred embodiment of this invention then looks at the next token in the line to determine if it is a maturity at **503** with the assumption that the current token is a coupon. In other words, the next token is used to provide more information about the current token. If there is no next token, it cannot be a maturity and the current token is determined not to be a coupon, and processing continues at **505**. If there is a next token, it is determined whether the next token is in another ticker's domain, and consequently whether it would apply to a new context instead of the current context. If the next token is in another ticker's domain, the current token is determined not to be a coupon, and processing continues at **505**. Also, if a maturity already exists in the current context, then it is determined that the next token is not a maturity and the current token is not a coupon. In this case, processing also continues at **505**. Further, if the next token is a number and a fraction, it is determined that it is not a maturity and that the current token is not a coupon. Processing then proceeds to **505**.

[0058] If after all of this analysis, the next token has not been resolved as a maturity, the next token is checked for compliance with a date format. If the next token is of the format MM/YY or YY or MM/YYYY, where Ms are numbers defining a month and Ys are numbers defining a year, or if the next token is a two digit integer preceded by a single quote, such as "04", then the next token is determined to have a date format. The next token may be preprocessed to remove day fields. For instance, a maturity of "12/5/04" can be preprocessed to be in the form 12/04. Although these particular formats are the preferred formats for a maturity date, one having ordinary skill in the relevant art will appreciate that the key point here is determining whether the next token has a date format. If the next token does not have a date format, it is determined not to be a maturity, and processing continues to **505**, the current token still being unresolved.

[0059] If the next token does have a date format, the next token is parsed to look for data that could not relate to a date, such as the number thirteen in a position where a month would be located, or a dollar sign. If it has any of these characteristics, the next token is determined not to be a maturity, and the current token not a coupon. Processing then proceeds to **505**. If the next token does not have any characteristic that would eliminate it from being a maturity, it is resolved as a maturity, and consequently, the current token is resolved as a coupon. Both the next token, resolved as a maturity, and the current token, resolved as a coupon, are stored in their respective fields in the current context at **504**. In this case, processing proceeds to **507** for an attempt to resolve a CUSIP for the current context.

[0060] Anytime a token has been resolved as a maturity, as just described, an attempt is made to resolve a CUSIP for the current context. The attempt to resolve a CUSIP is performed by comparing the data in the context at issue with the data in the bond list downloaded at 201 from the database 102. The ticker, coupon, and maturity data fields in the context at issue are compared with the CUSIPs in the bond list having the same ticker, coupon, and maturity. If the coupon field in the context at issue has a null value, all CUSIPs having the same ticker and maturity as the context are identified. If the maturity field in the context at issue has a null value, all CUSIPs having the same ticker and coupon as the context are identified. (This scenario could occur if no token in a line resolved as a maturity, at 402 in FIG. 4.) All identified CUSIPs are stored for later output, and may be stored in a data field in the context itself. In the case where one or more CUSIPs cannot be identified, processing proceeds normally, without any identified CUSIPs having been stored for later output. In the particular situation where an attempt to resolve or identify a CUSIP has been made after 507 in FIG. 5, processing continues on to the next token at 401 in FIG. 4.

[0061] Turning now to 505 in FIG. 5, if it was determined that the current token was not a coupon, the current token is then analyzed to determine if it is a maturity. If the current token is determined not to be a maturity, processing continues to 508. The manner in which the current token is determined to be or not be a maturity will now be described.

[0062] If a maturity exists in the current context, a decision is made that the current token cannot be a maturity. If the token is a number with a fraction, then it is determined not to be a maturity because date fields are not of this format. Also, the token must be able to resolve into a date format to be a maturity, and if it cannot, it is decided that it is not a maturity. As discussed with reference to 503, the preferred date formats are MM/YY or YY or MM/YYYY, with day fields having been preprocessed out of the token. If the current token does not have a date format, it is determined not to be a maturity, and processing continues to 508.

[0063] If the current token does have a date format, it is parsed to find data that could not relate to a date, such as the number 13 in a position where a month would be located, or a dollar sign. If it has any characteristic that would prevent it from being a date, the current token is determined not to be a maturity. If the current token does not have any characteristic that would eliminate it from being a maturity, it is determined to be a maturity. In this case, the token is stored as a maturity in the current context at 506. Also, since a token has been resolved as a maturity, an attempt is made to resolve a CUSIP for the current context at 507. After the attempt, the current token having been resolved as a maturity, processing of the next token begins at 401 in FIG. 4.

[0064] If the current token is not a coupon (501) or a maturity (505), it is determined whether it is a bid and/or offer at 508. A token that is to be a bid and/or an offer must have the following preferred formats: "N", "N/", "/N", or "N/N", where N represents a number. Whitespace can be before or after each N or "/", and each "P" can be replaced with a "-". Also, any tokens of this form that begin with a preceding zero are determined not to be bids and/or offers

because usually maturities begin with a zero. Examples of tokens that can be bids and/or offers are shown below in Table VIII.

TABLE VIII

100/
-90
60/65

[0065] In Table VIII, the "100 /" is a bid, the "-90" is an offer, and the "60/65" is an example of a token that includes both a bid and an offer, where the "60" is a bid and the "65" is an offer. Therefore, it is decided that the current token includes a bid if it is a number followed by a "/" or a "-", excluding whitespace. Also, if it is a number greater than or equal to 50 and is followed by the word "bid", it is determined to include a bid. Alternatively, it is determined that the current token includes an offer if it is a number preceded by a "/" or a "-", or if it is a number that is greater than or equal to fifty and is followed by the word "offer". A further optional way to help determine if the token includes a bid or an offer is to compare the number of total instances of the word "bid" or "offer" are present in the email with the number that have been processed.

[0066] If it is calculated that the current token includes a bid and/or an offer, the bid and/or offer data in the token is stored in the corresponding field(s) of the current context at 509. After storage, processing continues to the next token in the current email line at 401 in FIG. 4.

[0067] If it is calculated that the current token does not include a bid or an offer, the current token remains unresolved, and processing also continues to the next token in the current email line at 401 in FIG. 4.

[0068] Processing of the subsequent tokens in the line are the same as the process just described. Further, all of the tokens in the current line are processed, then each subsequent email line is processed (207 in FIG. 2), and when the email is completely processed (203 in FIG. 2), the stored security identifiers (CUSIPs) and their corresponding trade information, including bid price and/or offer price, are output at 204 in FIG. 2.

III. EXAMPLE

[0069] The processing depicted in FIGS. 3-5 will now be described with respect to an example. Suppose the line of an email shown in Table IX below is loaded for processing at 302 in FIG. 3.

TABLE IX

BAT	5.5	04/04	65-80	HHH	6.5	06	80-90
-----	-----	-------	-------	-----	-----	----	-------

[0070] At 303, it is determined that the line shown in Table IX is a data line because it contains at least the number 5.5, which could be a coupon, and the process then proceeds to 304 to prepare a list of tokens in this line. A token is considered to be a number or numbers separated by a "/" or a "-", and accordingly, the following tokens will be extracted from the line shown in Table IX: "5.5", "04/04", "65-80", "6.5", "06", and "80-90". The positions of each of these tokens in the line will also be recorded as "4", "8", "14", "24", "28", and "31", respectively, if the initial position in the line is considered to be zero.

[0071] At 305 in FIG. 3, which is elaborated upon in FIGS. 4 and 5, each of these tokens is processed as follows. At 401 in FIG. 4, it is determined that there are more tokens to process in this line because the six unprocessed tokens “5.5”, “04/04”, “65-80”, “6.5”, “06”, and “80-90” remain. At 403, the first token, “5.5” is selected. Because this is the initial token, and in the case of this example, it is assumed to be the initial token in the email, the initial ticker “BAT” is identified as a new ticker at 404. Because “BAT” is the initial ticker and “5.5” is the initial token, no previous context is determined to exist at 405, and a context is initialized for ticker “BAT” at 407. This context is initialized as shown in Table X below.

TABLE X

Context 1:	
Ticker:	BAT
Coupon:	Null
Maturity:	Null
Bid Price:	Null
Ask Price:	Null

[0072] At 501, the process of attempting to determine if the current token “5.5” is a coupon begins. First, the current context, context 1 shown in Table X, is checked to see if a coupon already exists in the context. Because the coupon field in context 1 has a value of “Null”, no coupon is determined to exist for this context and processing continues.

[0073] Next, it is determined if (1) the current token is less than or equal to the maximum coupon value in the list of bond data, (2) if a maturity exists in the current context, and (3) if the current token cannot be a bid or an offer, and if all three of these determinations are true, the current token is determined to be a coupon. However, since a maturity does not exist in context 1, this check fails and processing continues.

[0074] Next, it is determined whether the current token “5.5” is a number followed by a fraction or if it is preceded by a single quote, a “/”, a “-”, or a zero. If it is a number followed by a fraction or if it is preceded by a single quote, a “/”, a “-”, or a zero, it is determined not to be a coupon. However, “5.5” is not a number and a fraction, such as “5 1/2”, and it is not preceded by a single quote, a “/”, a “-”, or a zero, and processing continues.

[0075] Because the current token has not been resolved as a coupon as of yet, the next token “04/04” is checked to determine if it is a maturity at 503. But first, an inquiry is made as to whether the next token “04/04” is in a new ticker domain. However, since a new ticker is not between the position of the next token “04/04” and the position of the current token “5.5”, as shown in Table IX, it is decided that the next token is not in a new ticker domain. Further, because the maturity field in context 1 is “Null”, as shown in Table X, it is decided that a maturity for this context does not exist, and processing continues.

[0076] The next attempt to determine whether the next token “04/04” is a maturity includes checking it for compliance with a date format. Because “04/04” fits into a MM/YY format, where “M” represents a month digit and “Y” represents a year digit, and because “04/04” does not

have any characteristics that would prevent it from being a valid date, it is resolved as a maturity and the current token “5.5” is resolved as a coupon. Therefore, the current token “5.5” is stored as a coupon in the current context, context 1, and the next token “04/04” is stored in context 1 as a maturity at 504 in FIG. 5 and as shown in Table XI below.

TABLE XI

Context 1:	
Ticker:	BAT
Coupon:	5.5
Maturity:	04/04
Bid Price:	Null
Ask Price:	Null

[0077] At 507, an attempt to match one or more CUSIPs to the data in context 1 is made. That is, if any CUSIPs for ticker “BAT” with a coupon of “5.5” and a maturity of “04/04” exist, they will be identified and stored for later output. The CUSIP(s) that match the data in the current context may optionally be stored in the context itself. Whether or not one or more CUSIPs are identified, processing continues back to 401 in FIG. 4 to check for more unprocessed tokens.

[0078] The next unprocessed token is “65-80” as shown in Table IX, which is selected at 403. Since no new ticker is located between this token and the previous token, processing proceeds from 404 to 501 in FIG. 5, and the current token “65-80” is determined to be in the ticker domain of “BAT” and to apply to context 1

[0079] At 501 an attempt is made to resolve the current token “65-80” as a coupon. However, since a coupon already exists in context 1, as shown in Table XI, it is determined that the current token is not a coupon and processing proceeds to 505 to determine if it is a maturity. Similarly, because the current context includes a maturity, as shown in Table XI, the current token “65-80” is determined not to be a maturity and processing proceeds to 508 to check if it can be a bid and/or an offer.

[0080] At 508, the current token “65-80” is compared to the following bid/offer formats: “N”, “N/”, “/N”, or “N/N”, where N represents a number. Whitespace can be before or after each N or “/”, and each “/” can be replaced with a “-”. Also, bids and offers may not begin with a preceding zero. Because “65-80” has the format “N-N” and does not begin with a preceding zero, it is resolved as a bid and an offer and stored as such in the current context, context 1, as shown in Table XII below.

TABLE XII

Context 1:	
Ticker:	BAT
Coupon:	5.5
Maturity:	04/04
Bid Price:	65
Ask Price:	80

[0081] After storage of the bid and offer prices in context 1, processing continues back to 401 in FIG. 4 to find more unprocessed tokens in this line. The next unprocessed token

is "6.5" as shown in Table IX. At **403**, this token is selected as the current token, and the processing begins for determining what ticker domain this token belongs. To determine if the current token "6.5" is in a new ticker domain, a check is made for a ticker between the current token "6.5" and the previous token "65-80" at **403**. As shown in Table IX, the ticker "HHH" is between these tokens, and an answer of "yes" is returned at **404**. Context 1 now becomes the previous context at **405**, and another attempt to identify one or more CUSIPs for context 1 is made at **406**. After checking for CUSIPs at **406**, a new context, context 2 is initialized as shown in Table XIII below.

TABLE XIII

Context 2:	
Ticker:	HHH
Coupon:	Null
Maturity:	Null
Bid Price:	Null
Ask Price:	Null

[0082] The processing of the current token "6.5" and the remaining tokens "06" and "80-90" with respect to context 2 are processed in the same manner as the first three tokens were processed with respect to context 1 and will not be further described. Once processing of the email is complete, the CUSIPs identified for each context, if any, along with any resolved bid and/or offer prices pertaining to each context are output. According to experimental data, the invention extracts bond information from an assortment of emails having varying degrees of structure, 60% of the time, with 5-7% being false positives.

[0083] It is to be understood that the above-described embodiment and example is merely illustrative of the present invention and that many variations of the above-described embodiment and example can be devised by one skilled in the art without departing from the scope of the invention. For example, this system and method could easily be modified to scan partially unstructured documents for other information besides CUSIP numbers, and could be used, for instance, to scan email for SPAM, check files for viruses, or routing messages without specific addresses. It is therefore intended that any such variations and their equivalents be included within the scope of the following claims.

What is claimed is:

1. A method executed by a computer for processing partially unstructured data relating to a financial security, the method comprising:

- resolving first-identifying data from the partially unstructured data;
 - resolving second-identifying data from the partially unstructured data; and
 - determining whether a security is defined by the first-identifying data and the second-identifying data when the second-identifying data is of a predetermined type.
2. The method of claim 1,

wherein the first-identifying data represents a ticker,
 wherein the second-identifying data represents a coupon or a maturity, and
 wherein the predetermined type is a maturity.

3. The method of claim 1, further comprising:
 resolving third-identifying data from the partially unstructured data;
 determining whether a security is defined by the first-identifying data, the second-identifying data, and the third-identifying data.

4. The method of claim 3,
 wherein the first-identifying data represents a ticker,
 wherein the second-identifying data represents a coupon or a maturity,
 wherein the predetermined type is a maturity,
 wherein the third-identifying data represents a coupon if the second-identifying data represents a maturity, and
 wherein the third-identifying data represents a maturity if the second-identifying data represents a coupon.

5. The method of claim 3, further comprising:
 resolving trade information from the partially unstructured data.

6. The method of claim 5, further comprising:
 storing a security identifier representing a security determined to be defined by (a) the first-identifying data and the second-identifying data or (b) the first-identifying data, the second-identifying data, and the third-identifying data; and

outputting the security identifier and the trade information.

7. The method of claim 3,
 wherein said determining whether a security is defined by the first-identifying data and the second-identifying data determines whether a security is unambiguously defined by the first-identifying data and the second-identifying data, and

wherein said determining whether a security is defined by the first-identifying data, the second-identifying data, and the third-identifying data determines whether a security is unambiguously defined by the first-identifying data, the second-identifying data, and the third-identifying data.

8. The method of claim 5, further comprising:
 storing the first-identifying data in a context;
 storing the second-identifying data in the context;
 storing the third-identifying data in the context; and
 storing the trade information in the context.

9. A system for processing partially unstructured data relating to a financial security, the system comprising:

- an input device for receiving the partially unstructured data; and
- a processing device for performing actions comprising:
 - (a) resolving first-identifying data from the partially unstructured data;
 - (b) resolving second-identifying data from the partially unstructured data; and

(c) determining whether a security is defined by the first-identifying data and the second-identifying data when the second-identifying data is of a predetermined type.

10. The system of claim 9,

wherein the first-identifying data represents a ticker,

wherein the second-identifying data represents a coupon or a maturity, and

wherein the predetermined type is a maturity.

11. The system of claim 9, wherein the processing device performs actions further comprising:

resolving third-identifying data from the partially unstructured data; and

determining whether a security is defined by the first-identifying data, the second-identifying data, and the third-identifying data.

12. The system of claim 11,

wherein the first-identifying data represents a ticker,

wherein the second-identifying data represents a coupon or a maturity,

wherein the predetermined type is a maturity,

wherein the third-identifying data represents a coupon if the second-identifying data represents a maturity, and

wherein the third-identifying data represents a maturity if the second-identifying data represents a coupon.

13. The system of claim 11, wherein the processing device performs actions further comprising:

resolving trade information from the partially unstructured data.

14. The system of claim 13, further comprising:

a storage device; and

an output device,

wherein the processing device performs actions further comprising:

storing in the storage device a security identifier representing a security determined to be defined by (a) the first-identifying data and the second-identifying data or (b) the first-identifying data, the second-identifying data, and the third-identifying data; and

outputting, via the output device, the security identifier and the trade information.

15. The system of claim 11,

wherein said processing device determines whether a security is defined by the first-identifying data and the second-identifying data by determining whether a security is unambiguously defined by the first-identifying data and the second-identifying data, and

wherein said processing device determines whether a security is defined by the first-identifying data, the second-identifying data, and the third-identifying data by determining whether a security is unambiguously defined by the first-identifying data, the second-identifying data, and the third-identifying data.

16. The system of claim 13, further comprising:

a storage device,

wherein the processing device performs actions further comprising:

storing the first-identifying data in a context in the storage device;

storing the second-identifying data in the context in the storage device;

storing the third-identifying data in the context in the storage device; and

storing the trade information in the context in the storage device.

17. A method executed by a computer for identifying at least one of a plurality of predefined data vectors from partially unstructured data, said partially unstructured data comprising a plurality of data items having positions relative to each other in said partially unstructured data, the method comprising:

determining a position of each of one or more data items of a first type from the plurality of data items in the partially unstructured data;

selecting a data item of a second type from the plurality of data items in the partially unstructured data;

selecting one of the one or more data items of the first type based on its position relative to the selected data item of the second type; and

identifying a predefined data vector from the plurality of predefined data vectors using the selected data item of the first type and the selected data item of the second type.

18. The method of claim 17, further comprising:

selecting a data item of a third type from the plurality of data items in the partially unstructured data,

wherein identifying the predefined data vector using the selected data item of the first type and the selected data item of the second type identifies the predefined data vector using the selected data item of the first type, the selected data item of the second type, and the selected data item of the third type.

19. The method of claim 18, further comprising:

selecting a data item of a fourth type from the plurality of data items in the partially unstructured data.

20. The method of claim 19, further comprising:

outputting (a) an identifier representing the identified predefined data vector and (b) the data item of the fourth type.

21. The method of claim 20, further comprising:

storing in a context: (a) the selected data item of the first type, (b) the selected data item of the second type, (c) the selected data item of the third type, (d) the selected data item of the fourth type, and (e) the identifier.

22. The method of claim 17,

wherein the data item of the first type is a ticker, and

wherein the data item of the second type is a coupon or a maturity.

23. The method of claim 18,
wherein the data item of the first type is a ticker,
wherein the data item of the second type is a coupon or a
maturity, and
wherein the data item of the third type is a coupon or a
maturity, whichever the data item of the second type is
not.

24. The method of claim 19,
wherein the data item of the first type is a ticker,
wherein the data item of the second type is a coupon or a
maturity,
wherein the data item of the third type is a coupon or a
maturity, whichever the data item of the second type is
not, and

wherein the data item of the fourth type is trade informa-
tion.

25. The method of claim 20,
wherein the data item of the first type is a ticker,
wherein the data item of the second type is a coupon or a
maturity,

wherein the data item of the third type is a coupon or a
maturity, whichever the data item of the second type is
not,

wherein the data item of the fourth type is trade informa-
tion, and

wherein the identifier is a CUSIP.

* * * * *